arXiv:1609.09036v1 [q-bio.NC] 28 Sep 2016

High-Resolution Directed Human Connectomes and the Consensus Connectome Dynamics

Balázs Szalkai[a], Csaba Kerepesi[a,c], Bálint Varga[a], Vince Grolmusz[a,b,*]

[a] *PIT Bioinformatics Group, Eötvös University, H-1117 Budapest, Hungary*
[b] *Uratim Ltd., H-1118 Budapest, Hungary*
[c] *Institute of Computer Sceince and Control of the Hungarian Academy of Sciences*

Abstract

Here we show a method of directing the edges of the connectomes, prepared from diffusion tensor imaging (DTI) datasets from the human brain. Before the present work, no high-definition directed braingraphs (or connectomes) were published, because the tractography methods in use are not capable of assigning directions to the neural tracts discovered. Previous work on the functional connectomes applied low-resolution functional MRI-detected statistical causality for the assignment of directions of connectomes of typically several dozens of vertices. Our method is based on the phenomenon of the "Consensus Connectome Dynamics" (CCD), described earlier by our research group. In this contribution, we apply the method to the 423 braingraphs, each with 1015 vertices, computed from the public release of the Human Connectome Project, and we also made the directed connectomes publicly available at the site `http://braingraph.org`. We also show the robustness of our edge directing method in four independently chosen connectome datasets: we have found that 86% of the edges, which were present in all four datasets, get the very same directions in all datasets; therefore the direction method is robust, it does not depend on the particular choice of the dataset. We think that our present contribution opens up new possibilities in the analysis of the high-definition human connectome: from now on we can work with a robust assignment of directions of the connections of the human brain.

Introduction

Diffusion tensor imaging (DTI) datasets are widely used today for the construction of connectomes or braingraphs. The nodes (or vertices) of these graphs correspond to anatomically identified [1, 2], small (1-1.5 cm^2) areas of the gray

*Corresponding author
Email addresses: `szalkai@pitgroup.org` (Balázs Szalkai), `kerepesi@pitgroup.org` (Csaba Kerepesi), `balorkany@pitgroup.org` (Bálint Varga), `grolmusz@pitgroup.org` (Vince Grolmusz)

matter, and two nodes are connected by an undirected edge if neural fiber tracts are discovered, connecting the gray matter areas, corresponding to the two nodes [3, 4].

The considerable advantage of the graph approach to the analysis of the human brain is that the graphs retain the relevant anatomical connection information from the MRI data, but they do not contain the — mostly irrelevant — data on the exact orbits of the neural tracts in the white matter.

One of the most frequently used and largest diffusion MRI datasets are the public releases of the Human Connectome Project (HCP) [5]. Numerous publications are applying HCP data in different contexts, e.g., [6, 7, 8, 9, 10, 11, 12].

Directing the Connectome: fMRI-Based Methods

Diffusion tensor imaging methods cannot assign directions to the fiber tracts, and, consequently, to the edges of the connectomes mapped. Tractography methods can discover the orbits, that is, the geometric curves of the fiber tracts in the white matter, but — presently — cannot determine the correct one from the two possible directions of any given fiber tract. Consequently, the connectomes of braingraphs that are prepared from the diffusion MR imaging data are *undirected graphs*.

Functional MRI, on the other hand, is capable of mapping the temporal sequences of the activity of larger brain areas. Therefore, it seems to be possible to assign directions to these functional or effective connections between those larger areas of the gray matter. Several authors have attempted assigning directions to the edges of the connectomes using temporal sequences of activity changes, mostly on resting state fMRI or electro-physiology data, by applying a broad spectrum of statistical causality detection methods and models (e.g., [13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24]). The general approach of the articles listed can be described as follows: the temporal sequence of cerebral activity changes are translated to causality relations by statistical methods, and these causality relations are used to direct the connections between brain regions.

Since both the fMRI data and the non-invasive electro-physiological measurements have poor spatial resolution, usually the goal of these methods is not directing the edges of the diffusion tensor imaging-based, high-definition anatomical connectomes, but rather the functional connectomes on several dozens of vertices, corresponding to large areas of the gray matter, where haemodynamic activity changes were observed.

By the best of our knowledge, no significant attempts were made to direct the edges of structural connectomes (i.e., connectomes measured by DTI) with hundreds of vertices by these methods. One possible reason is due to the poor resolution of the fMRI: it is impossible to decide that if a causality is very probable between two cortical areas A and B in the direction from A to B, then which of the hundreds of the edges, running between A and B should be directed from A to B: theoretically, even one such connection may trigger the observed haemodynamic activity in B.

Consensus Connectomes

It is unusual in graph theory that hundreds of different graphs on the very same set of vertices are encountered. When we are working on braingraphs constructed from diffusion tensor imaging (DTI) data, this is the usual scenario: the cerebral areas of the different human subjects can be corresponded to the same reference brain map [2, 1], and, consequently, the graph edges in the connectome can be compared between subjects, since the vertices or nodes are named in each subject according to the very same anatomical brain map. This process is, in fact, a non-trivial refinement of the assignment of anatomical nomenclature to cerebral areas in printed brain anatomy atlases that appeared in the last several hundred years (e.g., [25, 26, 27]).

Consensus Connectome Dynamics (CCD) is a surprising phenomenon that was observed after our construction of the Budapest Reference Connectome Server [28, 29, 30]. For the description of the CCD phenomenon, we need to clarify some functions of the Budapest Reference Connectome Server.

The webserver at the address `http://connectome.pitgroup.org` is capable of finding and visualizing connectome edges that are present in more than a predefined number of connectomes. More exactly, let n denote the number of all connectomes processed by the server ($n = 6$ in version 1.0, $n = 96$ in version 2.0 and $n = 418$ or $n = 477$ — depending on other settings — in version 3.0). For any k: $0 \leq k \leq n$, we say that *the frequency of an edge* is k if the edge is present in exactly k connectomes out of the n ones [31]. Clearly, if we have n connectomes then the frequency of any edge is between 0 and n: it is 0 if the edge does not appear in any connectome at all; and it is n if it appears in all the n connectomes. The *k-consensus connectome* contains all the edges with a frequency greater than or equal to k: that is, each edge in the *k-consensus connectome* is present in at least k individual connectomes.

The Budapest Reference Connectome Server is capable of generating *k-consensus connectomes* in CSV and GraphML formats for further processing and independent visualization, and on the website users can also quickly view the resulting consensus connectomes.

Consensus Connectome Dynamics

After the publication of the article describing the Budapest Reference Connectome Server [28], a surprising property of the server was recognized [30]: when we generate k-consensus connectomes for $k = n$, next for $k = n - 1$,$k = n - 2$,..., and last for $k = 1$, then, naturally, more and more edges appear in the graphs (since the inclusion condition is weakened in every step). The surprising phenomenon is that the new edges do not appear randomly in the graph, but new edges are mostly connected to the already existing edges, forming a growing, tree-like structure as it is shown on the video `https://youtu.be/yxlyudPaVUE`, [30]. This "dynamical" observation is quantified on Fig. 2 of [30], and it is statically visualized on a very large component-tree at the address `http://pitgroup.org/static/graphmlviewer/index.html?src=connectome_dynamics_component_tree.g` (for the detailed — and not entirely obvious — description of the component-tree we refer to [30]).

In [30] we called this phenomenon "Consensus Connectome Dynamics", and hypothesized that the particular order of the appearance of the graph edges describes the order of growth of the axonal fibers of these connections in the individual brain development.

If the newly formed edges - i.e., fibers of axons - almost always are connected to the already developed edges - i.e., fibers of axons - then this hypothesis satisfyingly explains the Consensus Connectome Dynamics phenomenon, visualized at `https://youtu.be/yxlyudPaVUE`:

- Connections, represented by the graph edges that are present in most connectomes were grown first (e.g., edges in the n- or $n-1$-consensus connectomes).
- Next, the newer and newer edges were grown to be connected to the vertices of the older edges of the growing, tree-like structure, forming the k-consensus connectomes, for decreasing k values.

The newer and newer edges have gradually larger deviation, since small differences in the edge frequencies are added up in the growing structure: Clearly, if an edge e is not present in 20% of the graphs, then for another edge, say f, which connects to e, will hold that the e, f edge-pair cannot be present in at least 20% of the graphs, and most probably, the e, f pair will be missing from much more than 20% of all graphs, since f will be missing in some graphs, containing e. Consequently, the frequencies of the newer edges will be less than the frequencies of older edges. Therefore, we believe that the visualization at `https://youtu.be/yxlyudPaVUE`, describes *not only* the edges with gradually smaller frequencies, but also the order of growth of the connections in the human brain.

Results and Discussion

In the present work, we are assigning directions to some of the edges of braingraphs, depending on the appearance of that edge in the CCD.

The edge-directing method was first described in [30]; here we clarify the method, and make the 423 directed-edge connectomes publicly available at the site `http://braingraph.org`.

We have two requirements on the edge-directing method, based on Consensus Connectome Dynamics:

- **Feasibility:** While our direct knowledge is very limited on the axonal development of the human brain in general and the directions of fiber tracts in the connectome in particular ([32, 33]), we need to build onto the available biological observations in developing our method;

- **Robustness:** The directions, assigned to the edges of the connectomes need to be as independent as possible from the selected, specific sets of connectomes in the connectome sets in CCD trials. In other words, the CCD phenomenon needs the data of dozens or hundreds of connectomes; based on these connectomes and the CCD phenomenon, we assign directions to some of

the edges of the individual connectomes. However, we want these directions assigned to the edges of the individual braingraphs as independent as possible from the *particular* choice of the several dozen or several hundred connectomes included in the CCD probe.

Here we suggest an (a) Feasible and (b) Robust edge-directing method.

Feasibility

A large part of the axonal connections in the brain of mammals is developed in the embryonic state or the first several weeks after birth [33]. Microscopic studies of developing rat brains show that a considerable number of axons of the cortex (from layer V cortical neurons), in the early brain development, grow in the direction of subcortical regions [32], strengthening our CCD observation that neurons of the cortex are connected later to the subcortical structures than the connections develop between those subcortical structures.

Another important observation that forms a base of our method of directing the axonal fibers is the retrograde signaling of post-synaptic activation that stabilizes the newly formed synapses, and, consequently, prevents the retraction of the newly formed axon branches: work by numerous groups (e.g., [34, 35, 36, 37, 38, 39, 40]) witness this phenomenon.

Based on these observations, we make the following hypothesis: Neurons that are not connected to any other neuron yet, grow axons to reach other neurons.

(i) If the branch of the axon reaches a neuron that is not connected to other neurons yet, it retracts, since the post-synaptic activity of the un-networked neuron does not stabilize the connection (i.e., the synapse).

(ii) If the branch of the axon reaches an already "networked" neuron, that is, a neuron with connections to other neurons, then the post-synaptic activity of the "networked" neuron stabilizes the synapse and, consequently, the new connection.

Since the neuronal-level human braingraph with 80 billion neurons as vertices is unavailable today, we re-phrase the hypothesis above for the ROIs that form the 1015 vertices of our braingraphs, as follows:

ROIs that are not connected to any other ROI yet, grow axonal fibers to reach other ROIs as observed in the CCD phenomenon.

(iii) If the fiber reaches an ROI that is not connected to other ROIs yet, the fiber retracts, since the post-synaptic activities of the un-networked neurons do not stabilize the neuronal connections.

(iv) If the fiber reaches an already "networked" ROI, that is, an ROI with connections to other ROIs, then the post-synaptic activities of the "networked" neurons stabilize the synapses and the new connection will survive.

We stress that – instead of the mainly unknown order or temporal sequence of axonal development on the system level – we consider the phenomenon of the Consensus Connectome Dynamics, and we direct the edges of the connectomes according to the hypotheses (iii) and (iv) above.

Robustness

The robustness of the edge direction method is built into the edge direction algorithm, as it is detailed in the "Methods" section.

Methods

The braingraphs in this study were constructed from the data of the Human Connectome Project [5], using the workflow described in [28]. The undirected braingraphs have 1015 vertices each; the whole set of 423 braingraphs can be downloaded from the site `http://braingraph.org/download-pit-group-connectomes/`.

For verifying the robustness of the edge direction method, we first divided the braingraphs of the 423 subjects into four groups, the first contained 105, the others 106 graphs each. The graphs of each four groups were used – separately – to reproduce the CCD phenomenon, and based on the order of the appearance of the edges in the four different CCD probes, we assigned directions to some of the edges. Had the resulting directed graphs been radically different for the four groups, it would have indicated that our method is not a robust method of directing the edges of braingraphs.

As we detail below, it turned out that the four directed braingraphs were very similar to each other. After obtaining these four directed graphs, we used a majority-like approach to direct the edges of the consensus brain graph for the whole population, and, in a next step, the edges of the individual braingraphs.

Computing Consensus Braingraphs:

For each group, for $i = 1, 2, 3, 4$, we defined the consensus braingraph G_i as the multi-union of the individual connectomes. That is, for each possible graph edge we counted the subjects whose connectomes contained the specific edge. This number ranges from 0 to N_i, where $N_i = 106$, is the number of subjects in the given group. We defined the consensus braingraph as a simple graph containing all edges that were present in at least one of the subjects, where the edges are labeled (weighted) with the number of individual connectomes containing them.

Directing the edges with BFS:

We defined edge directions on this brain graph G_i as follows. For each k from N_i down to 1, we defined G_{ik} as the subgraph of G_i containing all edges with frequency at least k. This means all the edges which were present in at least k subjects within the ith group. Suppose that G_{ik} is already processed, and we want to process the graph $G_{i(k-1)}$, which is an augmentation of G_{ik} with

precisely those edges which are present in exactly $k-1$ individual graphs. Since G_{ik} is already processed, our task is now to direct these newly added edges only. To achieve this, we launch a BFS (breadth-first search) in $G_{i(k-1)}$, with G_{ik} as the source. This means that, for each node in $G_{i(k-1)}$, we calculate the distance of this node from G_{ik} within the graph $G_{i(k-1)}$. This distance is at least 0, is precisely 0 for the nodes of G_{ik}, is finite for the nodes connected to G_{ik} with a path of edges, and is infinite for the nodes of $G_{i(k-1)}$ not connected to G_{ik}. Nodes with infinite distance do exist because sometimes isolated edges or new, small graph components appear in $G_{i(k-1)}$.

After calculating the node distances, each edge in $G_{i(k-1)}$ will connect either two nodes of the same distance, or two nodes with a distance differing by exactly 1. If an edge connects u with v, and the distance of v is exactly 1 less than the distance of u, then we direct that edge from u to v. If an edge connects two equidistant nodes, then we leave this edge undirected and it will not acquire a direction in this consensus brain graph, not even in the subsequent steps.

This way, when we reach frequency $k = 1$ and have processed $G_{i1} = G_i$, we obtain a *partially directed* version of G_i. Some of its edges will become directed, and the remaining will be left undirected. As we do this for all four of the groups, we obtain four directed brain graphs on the same set of nodes.

Directed Consensus Connectome:

We unify these directed connectomes into one big directed consensus connectome as follows. If an edge is directed in the same direction in at least 2 of the 4 consensus graphs, and is undirected or directed in the other direction in at most 1 of the graphs, then the direction of this edge is defined as the most common value. For example, if an edge is directed forward in 2 graphs, backward in 1 graph, and is absent from the remaining graph, then the edge will become forward-directed in the result graph. On the other hand, if that edge is directed forward in 2 graphs, backward in 1 graph, and is undirected in the remaining graph, then its direction is considered ambiguous, and this edge will remain undirected in the result graph.

By running this algorithm, we could direct 48324 of the 71783 edges of the consensus brain graph for the whole population. This means that we could assign a direction to 67% of all the edges which were present in at least one of the braingraphs. We think that 67% as a very significant portion of the edges, since the edges include even those sporadic connections which were present in only one brain graph and thus cannot be oriented reliably. If we count only those edges which were present in all 4 consensus connectomes (there were 31873 such edges), then we see that 26305 of them were directed the same way in all 4 groups. This means that 82% of these 31873 edges acquired the same orientation in all 4 groups.

This proves the robustness of our method. Additionally, these numbers also show that the CCD phenomenon is robust, too.

How can we direct the individual braingraphs?

In the previous subsection we have presented a CCD-based method to direct most of the edges of the consensus connectome in a robust way. The edges of the individual braingraphs can be directed by applying the directions of the edges of the (partially) directed consensus braingraph as follows:

Let us consider an individual braingraph G and the directed consensus braingraph $\bar{G}$. If an edge e of G is directed in $\bar{G}$, then let us direct e in the same way in G, too. If edge e of G is undirected in $\bar{G}$, then e will remain to be undirected in G, too.

Data availability:

The Human Connectome Project's MRI data that were applied in the present work is available at: `http://www.humanconnectome.org/documentation/S500` [5]. The graphs (both undirected and directed) that were prepared by us from the HCP data can be downloaded at the site `http://braingraph.org/download-pit-group-connectomes/`.

The Budapest Reference Connectome Server [28, 29] is available at `http://connectome.pitgroup.org`. On that site, the reader can independently verify the phenomenon of the Consensus Connectome Dynamics by (i) choosing "Show options" (ii) and moving from right to left the "Minimum edge confidence" slider (iii) and observing the buildup of the edges on the visualization panel. If graphics problems appear then it is suggested to use another browser, e.g., Chrome.

The Consensus Connectome Dynamics is visualized on a video animation on `https://youtu.be/yxlyudPaVUE` for the whole brain and on `https://youtu.be/wBciB2eW6_8` for the frontal lobe only.

The individual braingraphs with directed edges can be accessed at the site `http://braingraph.org/download-pit-group-connectomes/`. Each graph refers to the original HCP ID in its filename.

Acknowledgments

Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

References

[1] Rahul S. Desikan, Florent Ségonne, Bruce Fischl, Brian T. Quinn, Bradford C. Dickerson, Deborah Blacker, Randy L. Buckner, Anders M. Dale,

R Paul Maguire, Bradley T. Hyman, Marilyn S. Albert, and Ronald J. Killiany. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *Neuroimage*, 31(3):968–980, Jul 2006. doi: 10.1016/j.neuroimage.2006.01.021. URL `http://dx.doi.org/10.1016/j.neuroimage.2006.01.021`.

[2] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

[3] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J. Honey, Van J. Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biol*, 6(7):e159, Jul 2008. doi: 10.1371/journal.pbio.0060159. URL `http://dx.doi.org/10.1371/journal.pbio.0060159`.

[4] Balázs Szalkai, Bálint Varga, and Vince Grolmusz. Graph theoretical analysis reveals: Women's brains are better connected than men's. *PLOS One*, 10(7):e0130045, July 2015. doi: doi:10.1371/journal.pone.0130045. URL `http://dx.plos.org/10.1371/journal.pone.0130045`.

[5] Jennifer A. McNab, Brian L. Edlow, Thomas Witzel, Susie Y. Huang, Himanshu Bhat, Keith Heberlein, Thorsten Feiweier, Kecheng Liu, Boris Keil, Julien Cohen-Adad, M Dylan Tisdall, Rebecca D. Folkerth, Hannah C. Kinney, and Lawrence L. Wald. The Human Connectome Project and beyond: initial applications of 300 mT/m gradients. *Neuroimage*, 80:234–245, Oct 2013. doi: 10.1016/j.neuroimage.2013.05.074. URL `http://dx.doi.org/10.1016/j.neuroimage.2013.05.074`.

[6] M. de la Iglesia-Vayá, J. Molina-Mateo, M. J. Escarti-Fabra, L. Martí-Bonmatí, M. Robles, T. Meneu, E. J. Aguilar, and J. Sanjuán. [magnetic resonance imaging postprocessing techniques in the study of brain connectivity]. *Radiologia*, 53(3):236–245, 2011. doi: 10.1016/j.rx.2010.11.007.

[7] Balázs Szalkai and Vince Grolmusz. Human sexual dimorphism of the relative cerebral area volumes in the data of the human connectome project. *arXiv preprint arXiv:1604.05992*, 2016.

[8] Matthew F. Glasser, Stamatios N. Sotiropoulos, J Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, Mark Jenkinson, and W. U-Minn H. C. P Consortium . The minimal preprocessing pipelines for the human connectome project. *Neuroimage*, 80: 105–124, Oct 2013.

[9] Balazs Szalkai, Balint Varga, and Vince Grolmusz. Mapping correlations of psychological and connectomical properties of the dataset of the human connectome project with the maximum spanning tree method. *arXiv:1602.04776*, feb 2016.

[10] Rick Herrick, Michael McKay, Timothy Olsen, William Horton, Mark Florida, Charles J. Moore, and Daniel S. Marcus. Data dictionary services in xnat and the human connectome project. *Front Neuroinform*, 8:65, 2014. doi: 10.3389/fninf.2014.00065. URL `http://dx.doi.org/10.3389/fninf.2014.00065`.

[11] Stephen M. Smith, Christian F. Beckmann, Jesper Andersson, Edward J. Auerbach, Janine Bijsterbosch, Gwenaëlle Douaud, Eugene Duff, David A. Feinberg, Ludovica Griffanti, Michael P. Harms, Michael Kelly, Timothy Laumann, Karla L. Miller, Steen Moeller, Steve Petersen, Jonathan Power, Gholamreza Salimi-Khorshidi, Abraham Z. Snyder, An T. Vu, Mark W. Woolrich, Junqian Xu, Essa Yacoub, Kamil Ugurbil, David C. Van Essen, Matthew F. Glasser, and W. U-Minn H. C. P Consortium . Resting-state fmri in the human connectome project. *Neuroimage*, 80:144–168, Oct 2013.

[12] Balázs Szalkai, Bálint Varga, and Vince Grolmusz. The graph of our mind. *arXiv preprint arXiv:1603.00904*, 2016.

[13] Alard Roebroeck, Elia Formisano, and Rainer Goebel. Mapping directed influence over the brain using granger causality and fmri. *Neuroimage*, 25 (1):230–242, 2005.

[14] Will D Penny, KE Stephan, A Mechelli, and KJ Friston. Comparing dynamic causal models. *NeuroImage*, 22(3):1157–1172, 2004.

[15] Sanggyun Kim, David Putrino, Soumya Ghosh, and Emery N. Brown. A granger causality measure for point process models of ensemble neural spiking activity. *PLoS Comput Biol*, 7(3): e1001110, Mar 2011. doi: 10.1371/journal.pcbi.1001110. URL `http://dx.doi.org/10.1371/journal.pcbi.1001110`.

[16] Emily SC Ching, Pik-Yin Lai, and CY Leung. Extracting connectivity from dynamics of networks with uniform bidirectional coupling. *Physical Review E*, 88(4):042817, 2013.

[17] Marc Timme and Jose Casadiego. Revealing networks from dynamics: an introduction. *Journal of Physics A: Mathematical and Theoretical*, 47(34): 343001, 2014.

[18] Björn Schelter, Matthias Winterhalder, Rainer Dahlhaus, Jürgen Kurths, and Jens Timmer. Partial phase synchronization for multivariate synchronizing systems. *Phys Rev Lett*, 96(20): 208103, May 2006. doi: 10.1103/PhysRevLett.96.208103. URL `http://dx.doi.org/10.1103/PhysRevLett.96.208103`.

[19] Jie Ren, Wen-Xu Wang, Baowen Li, and Ying-Cheng Lai. Noise bridges dynamical correlation and topology in coupled oscillator networks. *Phys Rev Lett*, 104(5):058701, Feb 2010. doi: 10.1103/PhysRevLett.104.058701. URL `http://dx.doi.org/10.1103/PhysRevLett.104.058701`.

[20] Olav Stetter, Demian Battaglia, Jordi Soriano, and Theo Geisel. Model-free reconstruction of excitatory neuronal connectivity from calcium imaging signals. *PLoS Comput Biol*, 8(8):e1002653, 2012. doi: 10.1371/journal.pcbi.1002653. URL http://dx.doi.org/10.1371/journal.pcbi.1002653.

[21] Felipe Gerhard, Tilman Kispersky, Gabrielle J. Gutierrez, Eve Marder, Mark Kramer, and Uri Eden. Successful reconstruction of a physiological circuit with known connectivity from spiking activity alone. *PLoS Comput Biol*, 9(7):e1003138, 2013. doi: 10.1371/journal.pcbi.1003138. URL http://dx.doi.org/10.1371/journal.pcbi.1003138.

[22] Yury V. Zaytsev, Abigail Morrison, and Moritz Deger. Reconstruction of recurrent synaptic connectivity of thousands of neurons from simulated spiking activity. *J Comput Neurosci*, 39 (1):77–103, Aug 2015. doi: 10.1007/s10827-015-0565-5. URL http://dx.doi.org/10.1007/s10827-015-0565-5.

[23] Stephen M. Smith, Karla L. Miller, Gholamreza Salimi-Khorshidi, Matthew Webster, Christian F. Beckmann, Thomas E. Nichols, Joseph D. Ramsey, and Mark W. Woolrich. Network modelling methods for fmri. *Neuroimage*, 54(2):875–891, Jan 2011. doi: 10.1016/j.neuroimage.2010.08.063. URL http://dx.doi.org/10.1016/j.neuroimage.2010.08.063.

[24] Jonathan W. Pillow, Yashar Ahmadian, and Liam Paninski. Model-based decoding, information estimation, and change-point detection techniques for multineuron spike trains. *Neural Comput*, 23(1):1–45, Jan 2011. doi: 10.1162/NECO_a_00058. URL http://dx.doi.org/10.1162/NECO_a_00058.

[25] Juan Valverde de Amusco. *Anatomia del corpo humano*. Per Ant. Salamanca, et Antonio Lafreri, 1560. URL https://www.nlm.nih.gov/exhibition/historicalanatomies/valverde_bio.html. NLM Unique ID: 2294026R.

[26] Sigismond Laskowski and Sigismond Balicki. *Anatomie normale du corps humain : atlas iconographique de XVI planches / par le docteur S. Laskowski ; dessinee's d'apres les preeparations de l'auteur par S. Balicki.* Geneve : Braun, 1894. URL https://www.nlm.nih.gov/exhibition/historicalanatomies/laskowski_home.html.

[27] Ferenc Kiss and Janos Szentagothai. *Atlas of Human Anatomy*. Springer, 1987.

[28] Balázs Szalkai, Csaba Kerepesi, Bálint Varga, and Vince Grolmusz. The Budapest Reference Connectome Server v2. 0. *Neuroscience Letters*, 595: 60–62, 2015.

[29] Balázs Szalkai, Csaba Kerepesi, Bálint Varga, and Vince Grolmusz. Parameterizable consensus connectomes from the human connectome project: The budapest reference connectome server v3.0. *arXiv preprint arXiv:1602.04776*, feb 2016.

[30] Csaba Kerepesi, Balázs Szalkai, Bálint Varga, and Vince Grolmusz. How to direct the edges of the connectomes: Dynamics of the consensus connectomes and the development of the connections in the human brain. *PLOS One*, 11(6):e0158680, June 2016. URL http://dx.doi.org/10.1371/journal.pone.0158680.

[31] Csaba Kerepesi, Balint Varga, Balazs Szalkai, and Vince Grolmusz. The dorsal striatum and the dynamics of the consensus connectomes in the frontal lobe of the human brain. *arXiv*, 1605.01441, May 2016.

[32] J. A. De Carlos and D. D. O'Leary. Growth and targeting of subplate axons and establishment of major cortical pathways. *J Neurosci*, 12(4): 1194–1211, Apr 1992.

[33] Tommy L Lewis, Jr, Julien Courchet, and Franck Polleux. Cell biology in neuroscience: Cellular and molecular mechanisms underlying axon formation, growth, and branching. *J Cell Biol*, 202(6):837–848, Sep 2013. doi: 10.1083/jcb.201305098. URL http://dx.doi.org/10.1083/jcb.201305098.

[34] Bing Hu, Angeliki Maria Nikolakopoulou, and Susana Cohen-Cory. Bdnf stabilizes synapses and maintains the structural complexity of optic axons in vivo. *Development*, 132(19):4285–4298, Oct 2005. doi: 10.1242/dev.02017. URL http://dx.doi.org/10.1242/dev.02017.

[35] H. T. Cline and M. Constantine-Paton. Nmda receptor antagonists disrupt the retinotectal topographic map. *Neuron*, 3(4):413–426, Oct 1989.

[36] I. Rajan, S. Witte, and H. T. Cline. Nmda receptor activity stabilizes presynaptic retinotectal axons and postsynaptic optic tectal cell dendrites in vivo. *J Neurobiol*, 38(3):357–368, Feb 1999.

[37] J. T. Schmidt, M. Buzzard, R. Borress, and S. Dhillon. Mk801 increases retinotectal arbor size in developing zebrafish without affecting kinetics of branch elimination and addition. *J Neurobiol*, 42(3):303–314, Feb 2000.

[38] Edward S. Ruthazer and Hollis T. Cline. Insights into activity-dependent map formation from the retinotectal system: a middle-of-the-brain perspective. *J Neurobiol*, 59(1):134–146, Apr 2004. doi: 10.1002/neu.10344. URL http://dx.doi.org/10.1002/neu.10344.

[39] D. J. Zou and H. T. Cline. Control of retinotectal axon arbor growth by postsynaptic camkii. *Prog Brain Res*, 108:303–312, 1996.

[40] John T. Schmidt. Activity-driven sharpening of the retinotectal projection: the search for retrograde synaptic signaling pathways. *J Neurobiol*, 59(1):114–133, Apr 2004. doi: 10.1002/neu.10343. URL http://dx.doi.org/10.1002/neu.10343.